Volume Title: Analyses in the Economics of Aging

Volume Author/Editor: David A. Wise, editor

Volume Publisher: University of Chicago Press

Volume ISBN: 0-226-90286-2

Volume URL: http://www.nber.org/books/wise05-1

Publication Date: August 2005

Title: Passive Decisions and Potent Defaults

Author: James J. Choi, David Laibson, Brigitte C. Madrian, Andrew Metrick

URL: http://www.nber.org/chapters/c10357

2

Passive Decisions and Potent Defaults

James J. Choi, David Laibson, Brigitte C. Madrian, and Andrew Metrick

Default options have an enormous impact on household choices. Such effects have now been extensively documented in the literature on 401(k) plans (Madrian and Shea 2001; Choi et al. 2002b, 2004). Defaults have been shown to affect participation, savings rates, rollovers, and asset allocation. For example, Choi et al. (2004) study three firms that use automatic enrollment. When employees at these firms are automatically enrolled in their 401(k) plan, only a tiny fraction opt out, producing participation rates exceeding 85 percent regardless of tenure. But when employees at these firms were not automatically enrolled, participation rates were significantly lower, ranging from 26 percent to 43 percent after six months of tenure, and from 57 percent to 69 percent after three years of tenure.

James J. Choi is a Ph.D. candidate in economics at Harvard University. David Laibson is a professor of economics at Harvard University and a research associate of the National Bureau of Economic Research. Brigitte C. Madrian is the Boettner Associate Professor in Financial Gerontology and an associate professor of business and public policy at the Wharton School, University of Pennsylvania, and a research associate of the National Bureau of Economic Research. Andrew Metrick is an associate professor of finance at the Wharton School, University of Pennsylvania, and a faculty research fellow of the National Bureau of Economic Research.

This paper is a longer version of a paper titled "Optimal Defaults" (Choi et al. 2003). We thank Hewitt Associates for their help in providing the data. We are particularly grateful to Lori Lucas, Jim McGhee, and Scott Peterson, three of our many contacts at Hewitt. We are grateful for the comments of Robert Barro, Robert Hall, Mark Iwry, Antonio Rangel, Richard Thaler, and seminar participants at Harvard, Massachusetts Institute of Technology, the National Bureau of Economic Research, the University of Chicago, the University of Minnesota, the University of Southern California, and Wharton. Numerous research assistants have contributed to this work. We are particularly indebted to Nelson Uhan. We acknowledge financial support from the National Institute of Aging (grant R01AG021650). Choi acknowledges support from a National Science Foundation graduate research fellowship. Laibson acknowledges financial support from the National Science Foundation (grant 0099025).

Defaults matter for three key reasons that we model in this paper. First, acts of commission—for example, opting out of a default—are costly. Second, these costs change randomly over time and therefore generate an option value of waiting to change a default. Decision makers would like to wait for a low-cost period (e.g., a free weekend) to make a change. Third, people have a tendency to procrastinate. Even if they want to make a change, they have a tendency to delay that change longer than they should.

Because of these effects, the choice of a particular default can have a significant effect on consumer welfare. However, it is not always obvious how to select a socially optimal default.

If *all* employees would like to be saving at a rate of exactly 5 percent in their 401(k) plan, then the employees' welfare will be maximized if the employer sets a 5 percent default. But the calculation of an optimal default is not as straightforward if different employees have different optimal savings rates. For example, what is the optimal default savings rate if employees have optimal savings rates that are distributed uniformly with a mean of 5 percent?

In this paper, we develop a theory of optimal defaults that implies that the obvious answer to the previous question—5 percent—is not necessarily the right answer. In a world of heterogeneous agents, it may sometimes be optimal to set *extreme* defaults that are far away from the mean optimal savings rate. This effect arises for two reasons. First, a default that is far away from a consumer's optimal savings rate may make that consumer *better off* than a default that is closer to the consumer's optimal savings rate. Intuitively, if an agent suffers from a procrastination problem, then a "bad" default—that is, one that is far from the consumer's optimal savings rate—will be more motivating than a better default. Hence, sometimes bad defaults make people better off than better but imperfect defaults. Second, our theory implies that optimal defaults are highly sensitive to the actual distribution of optimal savings rates. In particular, optimal defaults are often associated with the modal optimal savings rate and not the mean optimal savings rate. Since these modes are sometimes extreme (e.g., minimum or maximum contribution rates), optimal defaults will sometimes be extreme as well.

At the end of our paper we calibrate our model and use it to calculate optimal defaults for employees at four different companies. For two of our companies, the optimal default is close to the mean optimal savings rate, whereas for the other two companies the optimal defaults are extreme: 0 percent and 15 percent, respectively. Our work suggests that optimal defaults are likely to be at one of three savings rates: the minimum savings rate (0 percent), the employer match threshold (typically 5 percent or 6 percent), or the maximal savings rate (around 15 percent in our sample of companies from the late 1990s).[1]

1. More recently, regulatory changes under the Economic Growth and Tax Relief Reconciliation Act of 2001 (EGTRRA) have led many companies to raise their maximum savings rates well above the historical norm of 15 percent.

2.1 A Model of Savings Choices

We adapt the model of Choi et al. (2002a) to describe the 401(k) enrollment decisions of employees that have been newly hired at a firm. However, the model is general enough to describe any problem in which an actor decides when to move from a default state s_D to an optimal state s^*.

We assume that each employee at a firm has a fixed optimal savings rate (i.e., optimal state) s^*, with density function f characterizing the distribution of these optimal savings rates for the population of employees in the firm. When new employees join the firm, the employees are automatically enrolled at a default savings rate of s_D, which is a choice variable for the firm. In this paper, we consider the case in which this default can only take values in the support of f.[2] We assume that the firm uses a single default savings rate for all of its employees either because the firm does not observe an employee's true type, s^*, or because of legal or practical costs of implementing employee-specific defaults.[3]

Employees remain at the default election s_D unless they opt out of the default by incurring a cost c. This opt-out cost is drawn each period and takes the value 1 with probability $\mu > 1$ and value 0 with probability $1 - \mu$. The value of the cost is known when the agent decides on her action. We suppress individual and time subscripts to simplify notation.

When the agent opts out, she sets her savings rate equal to her optimal savings rate s^*, which we assume the agent knows with certainty.[4] Until that action takes place, the agent suffers a flow loss of $L = L(s_D, s^*) \geq 0$, where the first argument of L is the current savings rate and the second argument of L is the optimal savings rate. After the action occurs, the agent suffers a flow loss of $0 = L(s^*, s^*)$.

Finally, we assume that agents are *naive* hyperbolic discounters, with discount function $1, \beta\delta, \beta\delta^2, \ldots$.[5] Such naive agents believe that their future selves will make choices that are consistent with their current preferences. We adopt such naive beliefs because they increase the force of procrastination, but our qualitative results would be unchanged if we instead assumed that agents are sophisticated in their beliefs. For sim-

2. See Choi et al. (2002a) for a generalization.

3. Such employee-specific defaults are a natural extension of our current framework and merit theoretical and practical evaluation.

4. Another natural generalization is to consider the case in which agents have imperfect information about their personal value of s^*. If agents learn more about this value over time, they have another motive for delaying the costly action of opting out of the default.

5. See Laibson (1997) for a discussion of hyperbolic discount functions and Akerlof (1991) and O'Donoghue and Rabin (1999) for a discussion of naifs and procrastination. Note that the term "hyperbolic" is overly restrictive, since the important property of these preferences is simply that they are characterized by more discounting in the short run than in the long run.

plicity and analytical tractability, we set $\delta = 1$ (no long-run discounting).[6] We also adopt the standard hyperbolic assumption of $\beta < 1$.

We use the following timing convention. If the employee has not previously opted out of the default, the period begins with a flow loss of L. The employee then draws a current opt-out cost c and decides whether to delay opting out or to instead pay the cost, thereby ending the game. If the employee delays, she will pay a flow cost of L next period and also face an anticipated continuation value function, which we denote $v(c')$, where c' represents next period's draw from the cost distribution. Hence, the employee chooses to pay c and end the game if the cost today is less than the discounted cost of delay, or

$$c < \beta\,[L + Ev(c')].$$

When this inequality is not satisfied, the employee chooses to delay. Ignoring mixed strategies, which only arise on a zero measure region of the parameter space, the employee's strategy is thus

$$\text{(1)}\qquad \begin{array}{ll} \text{“Opt out only when } c = 0\text{”} & \text{if } \beta\,[L + Ev(c')] < 1 \\ \text{“Opt out only when } c = 0 \text{ or } c = 1\text{”} & \text{if } \beta\,[L + Ev(c')] \geq 1 \end{array}$$

2.1.1 Naive Expectations and the Continuation Value Function $v(c)$

Since the employee is assumed to be a naive hyperbolic agent, the continuation value function is constructed under the (mistaken) belief that all future selves will exhibit no time discounting, since this is what today's self wants those future selves to do. Recall that $\delta = 1$.

The strategy of opting out whatever the draw from the cost distribution means that the employee's expected loss is $\mu = E(c)$. Waiting until $c = 0$ to opt out implies that the employee's expected loss would be

$$\begin{aligned} Ev(c \mid \text{wait until } c = 0) &= \mu[L + Ev(c \mid \text{wait until } c = 0)] \\ &= \frac{\mu L}{1 - \mu}. \end{aligned}$$

This formula has a natural interpretation: the expected costs are equal to the expected per-period loss, μL, multiplied by the expected duration of the losses, $1/(1 - \mu)$.

If $L < 1 - \mu$, then $\mu L/(1 - \mu) < \mu$, implying that the expected losses generated by waiting to opt out until $c = 0$ are less than the losses from opting out immediately at cost $c = 1$. So if $L < 1 - \mu$, the employee will plan to wait until $c = 0$ to opt out. If $L \geq 1 - \mu$, the employee anticipates that next period she will act with certainty. In summary,

6. We will calibrate our model at the frequency of a pay cycle. So if the annual long-run discount rate is 0.05, then the discount rate per pay cycle is approximately $0.05/26 = 0.002$ or $0.05/12 = 0.004$, implying respective δ values of 0.998 and 0.996. Relative to these values, setting $\delta = 1$ has little impact on our results.

$$(2) \qquad Ev(c) = \begin{cases} \dfrac{\mu L}{1-\mu} & \text{if } L < 1-\mu \\ \mu & \text{if } L \geq 1-\mu \end{cases}$$

We reiterate that $Ev(c)$ is based on naive beliefs, so this expectation reflects the actor's incorrect model of her future behavior.

2.1.2 Actual Actions and Welfare

Using equations (1) and (2), the probability of opting out in any period will be

$$p = \begin{cases} 1-\mu & \text{if } L < \dfrac{1}{\beta} - \mu \\ 1 & \text{if } \dfrac{1}{\beta} - \mu \leq L \end{cases}$$

So the expected cost of opting out, conditional on opting out, will be

$$E(c \mid \text{opt out}) = \begin{cases} 0 & \text{if } L < \dfrac{1}{\beta} - \mu \\ \mu & \text{if } \dfrac{1}{\beta} - \mu \leq L \end{cases}$$

Let $w(c)$ represent the employee's expected total costs, discounted with the agent's long-run discount factor. A recursive representation for $w(c)$ is given by

$$\begin{aligned} Ew(c) &= pE(c \mid \text{opt out}) + (1-p)\delta[L + Ew(c')] \\ &= pE(c \mid \text{opt out}) + (1-p)[L + Ew(c')]. \end{aligned}$$

We evaluate social welfare using the long-run discount factor δ and omitting the short-run discount factor β. These preferences represent the actor's preferences at economic birth, which we assume occurs before she starts working at the firm. The last equation contains *no discounting,* since it reflects the fact that $\delta = 1$ in our calibration. Note however that our results would not change qualitatively if we had instead assumed $\delta < 1$ throughout our analysis.

Because $Ew(c) = Ew(c')$, we can show that

$$Ew(c) = \begin{cases} \dfrac{\mu L}{1-\mu} & \text{if } L < \dfrac{1}{\beta} - \mu \\ \mu & \text{if } \dfrac{1}{\beta} - \mu \leq L \end{cases}$$

We are now in a position to characterize the relationship between defaults and welfare. To do this, we consider the relationship between expected (dis)utility and L, the per-period flow losses of not being at an optimum. To focus on the role of L, we stop suppressing L in our notation and consider

$$W(L) = Ew(c)_{|L}.$$

$W(L)$ is the expected losses for an agent with initial flow losses per period of L.

In a standard model with exponential discounting (i.e., $\beta = 1$), $W(L)$ would increase as flow costs L increase. But for hyperbolics (i.e., $\beta < 1$), it will always be the case that W is nonmonotonic in L. To see this, note that $W(L) = \mu$ when $L = 1 - \mu$. This is the level of L at which an exponential (i.e., dynamically consistent) agent should opt out of the default whatever the cost realization. But when $c = 1$, a hyperbolic agent will only opt out of the default if $L \geq (1/\beta) - \mu$, which is greater than $1 - \mu$. Hence, when $1 - \mu < L < (1/\beta) - \mu$, the hyperbolic agent is insufficiently motivated to act, and this motivational gap produces self-defeating procrastination. In this region of L values, the expected loss function lies above μ, the value that $W(L)$ would take if the agent were not procrastinating and were willing to act at the high cost realization. But once L is high enough—specifically, above $(1/\beta) - \mu$—the procrastination effect vanishes and expected costs fall back to μ, since the hyperbolic agent is now willing to act whatever the cost realization. Figure 2.1 plots the expected cost function against the flow costs L, revealing the nonmonotonicity that arises whenever $\beta < 1$.

In a world with procrastination, moving the agent further from the optimum (i.e., increasing flow costs L) can make an agent *better* off, since it decreases the agent's tendency to procrastinate. This effect is not everywhere offset by the direct effect of reduced welfare arising from the increase in the delay cost, L.

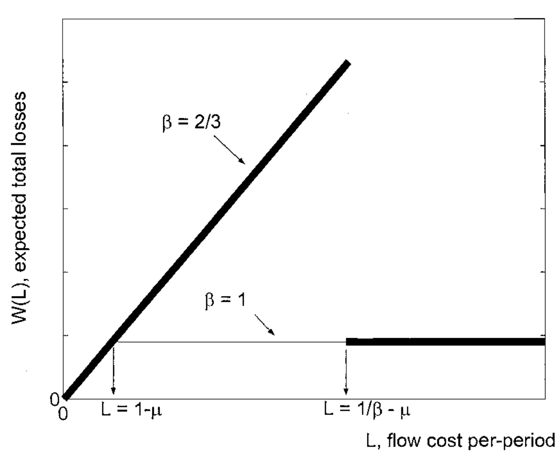

Fig. 2.1 Expected total losses as a function of flow cost per period

2.1.3 The Firm's Optimization Problem

We now analyze the employer's choice of a default savings rate under the assumption that the employer is interested in maximizing the welfare of the firm's employees. We recognize, however, that employer and employee incentives need not generally be aligned. This is particularly likely in the case presented here, since naive hyperbolic agents will not anticipate their own tendency to procrastinate and hence will not pick an employer based on the employer's ability to mitigate the harms of such procrastination. Therefore, this normative exercise is also relevant for regulators or unions that can influence the defaults that firms pick. Identifying and incorporating the other motivations and constraints that firms face in designing their benefit plans (e.g., nondiscrimination testing, good corporate citizenship, reputational value in the labor market, or personal altruism, to name a few) is beyond the scope of the current paper.

We derive the optimal default, s_D^*, that minimizes the social welfare function,

$$\int_{\underline{s}}^{\bar{s}} W(L(s_D, s^*))f(s^*)ds^*. \tag{3}$$

We adopt the cost function

$$L(s_D, s^*) = \kappa(s_D - s^*)^2.$$

This quadratic cost function is convex in deviations from the optimal savings rate, s^*, and has the advantage of analytic tractability. However, it does not reflect the particular institutional features of many 401(k) plans (e.g., an employer match that ends at a threshold, implying a discontinuity in the cost function). We believe that the quadratic cost function represents a good compromise between tractability and realism.

We will minimize equation (3) numerically, using the actual estimated distribution of optimal savings rates. However, for the purposes of exposition, it is useful to consider the case in which $f(s^*)$ is uniform over support $[\underline{s}, \bar{s}]$. In this case, one can prove the following result when $\beta < 1$:

$$s_D^* = \begin{cases} \dfrac{\bar{s} + \underline{s}}{2} & \text{if } \bar{s} - \underline{s} \text{ is small} \\ \underline{s} + \sqrt{\dfrac{1}{\kappa}(1-\mu)} \text{ or } \bar{s} - \sqrt{\dfrac{1}{\kappa}(1-\mu)} & \text{if } \bar{s} - \underline{s} \text{ is large} \end{cases}.$$

Intuitively, when there is little variation in optimal savings rates, it is best to design a default that is in the middle of the range of optimal savings rates, since all employees will then be very close to their optimal savings rate and delays in opting out of the default will not be very costly. By contrast, when there is a great deal of variation in optimal savings rates, it is better to design a default that is close to one of the two boundaries of the support. This "boundary" strategy reduces the proportion of employees who engage in costly procrastination, since the boundary strategy reduces the fraction of employees who fall in the "procrastination" interval $1 - \mu < L < (1/\beta) - \mu$.

Finally, note that if $\beta = 1$ and f is uniform, then $s_D^* = (\bar{s} + \underline{s})/2$ will always be an optimum[7] because the procrastination effect does not apply and there is no gain in welfare from moving agents away from their optima.

It is also useful to emphasize a trivial property of these models, which is important in the empirical analysis that follows. This additional effect is easiest to understand if we assume that f is a discrete density on the domain of feasible savings rates: $\{0.00, 0.01, 0.02, \ldots\}$. Then it is easy to show that

$$\lim_{\kappa \to \infty} s_D^* \in \arg\max_{s^*} f(s^*).$$

In other words, as the cost of deviations rises ($\kappa \to \infty$), the optimal default converges to the mode of the distribution of s^*. This effect is driven by the fact that for large costs of deviating from s^*, all employees will immediately adjust to their s^* except those who are already at their optima. Hence, the optimal social policy minimizes adjustment costs by setting the default equal to the most common value of s^*. We refer to this as the mode effect.

2.1.4 Calibration

Our model has very few free parameters: the density of optimal savings rates, $f(s^*)$; the short-term discount factor, β; the scaling variable, κ; and the probability of a high-cost draw, μ. We further restrict this list by using individual employee data to pin down the density f (see next section). We set $\beta = 2/3$ reflecting a large body of experimental evidence and a growing body of field evidence. For example, Laibson, Repetto, and Tobacman (2003) use the method of simulated moments to estimate β using household financial data. Their benchmark estimate is 0.70 with a standard error of 0.11.

Only κ and μ remain to be calibrated. Before doing this we need to pick units for the variables in our model. We assume that time units are periods of a pay cycle (about two weeks). We assume that utility units can be interpreted in terms of a money metric in which one unit of utility is equal in value to one-tenth of a pay cycle of income. So when the cost realization is high ($c = 1$), opting out of the default generates a time cost that is equal in value to one-tenth of the agent's income during that pay cycle. We assume that such busyness is the norm and set $\mu = 0.9$. It then follows that the cost realization will be zero $0.1 = (1 - \mu)$ of the time.

7. However, it will not generally be the unique optimum.

To set κ, we use the following thought experiment. Suppose that a consumer is 10 percentage points away from her optimal savings rate: $|s_D - s^*| = 0.1$. What is the money-metric cost of this deviation? Let x represent the loss in units of one-tenth of one pay cycle of income. Then, $\kappa(0.1)^2 = x$. We will consider a range of values for x: 0.1, 1, 10. This translates into the following range of values for κ: 10, 100, 1000. We consider this wide range for two reasons. First, we are agnostic about the appropriate calibration value. Second, we wish to explore the sensitivity of our results to the choice of κ. However, if forced to choose, we would set $\kappa = 100$, implying that a 10 percentage point deviation in one's savings rate is as bad as losing one-tenth of one's income during that pay cycle. For companies with an employer match, one could motivate losses of this magnitude by considering the missed match payments induced by undersaving.

2.2 Empirical Analysis

Table 2.1 shows the variation in both 401(k) plan design and employee characteristics of the four companies for which we compute the optimal default 401(k) savings rate. We denote these four companies by their industry: Health, Office, Food, and Finance. All are large employers with well-established 401(k) plans.

There are two key differences in the 401(k) plan environment that vary across the companies. First, two of the companies (Health and Office) match employee contributions up to 6 percent of pay, while the other two have no match at all. These latter companies are of interest because the distribution of employee contribution rates will not be affected by the pres-

Table 2.1 **Characteristics of employees and their 401(k) plans**

	Health company (1)	Office company (2)	Food company (3)	Financial company (4)
Employer match	\$0.50/\$1 to 6%	\$0.67/\$1 to 6%	None	None
Contribution rate range (%)	0 to 15	0 to 16	0 to 15	0 to 15
Company DB plan	No	Yes	Yes	No
401(k) participation rate (%)	61.9	74.2	32.8	63.4
Avg. 401(k) contribution rate (%)	4.3	4.5	2.0	6.0
Median salary (\$)	31,034	27,629	25,355	41,109
Median age (years)	37.9	36.7	38.5	28.9
Median tenure (years)	4.8	5.4	5.6	2.0
Fraction female (%)	77.8	30.1	54.0	50.0
Year	1997	1998	1998	1998

Source: Company summary plan descriptions and calculations of the authors.

Notes: The sample in column (1) is all employees with 1+ year of tenure. The sample in column (2) is all employees with 2+ years of tenure. The sample in columns (3) and (4) is all employees.

ence of a match threshold. Having an employer match may either raise or lower the desired 401(k) contribution rate. Because the match subsidizes saving in the 401(k) plan, employees with a match may desire to contribute more, at least up to the match threshold. However, the match also increases the total amount of savings that is being done, and the employees may use the match as a means to offset their own contributions.

The second key difference in plan environment is that two of the companies (Office and Food) have an employer-sponsored defined benefit pension plan, while the other two do not. Other things being equal, we would expect a lower desired savings rate for employees in companies with a defined benefit pension.

The workforce demographics of our four companies also vary quite considerably. The median pay ranges from \$25,000 per year in Food to \$41,000 per year in Finance. Because Social Security replaces a higher fraction of income for low-income employees, we would expect a higher desired savings rate for high-income employees. There is also significant variation in the fraction of employees that are female (from 30 percent in Office to 78 percent in Health) and the median age of the workforce (from twenty-nine years in Finance to thirty-nine years in Food).

To estimate the distribution of optimal savings rates (i.e., the density f in the model), we use two approaches. First, we report densities over 401(k) savings rates for "medium-tenure" employees. We informally reason that such medium-tenure employees have been at a firm long enough to select their optimal savings rate (i.e., the option value of waiting and procrastination hurdles have been surmounted) but not so long that tenure-driven selection effects dominate the data. These savings densities are reported in table 2.2 for employees with three to five years of tenure (density f_1) and five to seven years of tenure (density f_2).

Second, we use a regression framework to control for demographic variables. We run an ordered logit regression in which the explanatory variable is the actual 401(k) contribution rate chosen by each individual employee. We include nonparticipation, which implies a 0 percent contribution rate, as one of the categories. The control variables in the regressions are ln(pay), ln(age), ln(tenure), and a gender dummy variable ($D = 1$ if the employee is female). We then predict the distribution of contribution rates that would obtain if each employee had thirty years of tenure, holding other demographic characteristics constant. The underlying presumption behind this exercise is that thirty years is enough time to overcome any delays due to procrastination or the option value of waiting. The projected density from this procedure is reported as density f_3 in table 2.2.

With these densities in hand, we are now in a position to estimate the optimal savings rate by minimizing equation (3), the social welfare function. We undertake this minimization for $3 \times 3 \times 4$ cases of interest: three different values for κ, three different ways of calculating the density f, and four

Table 2.2 **Savings rate distributions**

	Health company			Office company			Food company			Financial company		
s^*(%)	$f_1(s^*)$	$f_2(s^*)$	$f_3(s^*)$	$f_1(s^*)$	$f_2(s^*)$	$f_3(s^*)$	$f_1(s^*)$	$f_2(s^*)$	$f_3(s^*)$	$f_1(s^*)$	$f_2(s^*)$	$f_3(s^*)$
0	0.35	0.26	0.20	0.32	0.23	0.12	0.72	0.59	0.51	0.22	0.19	0.19
1	0.01	0.01	0.01	0.01	0.01	0.01	0.02	0.04	0.04	0.02	0.02	0.01
2	0.06	0.06	0.03	0.15	0.14	0.09	0.04	0.06	0.06	0.02	0.02	0.02
3	0.06	0.07	0.05	0.05	0.05	0.04	0.04	0.05	0.05	0.04	0.02	0.02
4	0.04	0.05	0.03	0.04	0.05	0.04	0.02	0.03	0.03	0.03	0.03	0.02
5	0.04	0.05	0.04	0.04	0.04	0.04	0.06	0.09	0.12	0.09	0.09	0.07
6	0.26	0.28	0.30	0.24	0.31	0.36	0.01	0.02	0.02	0.04	0.05	0.04
7	0.02	0.02	0.02	0.01	0.01	0.02	0.01	0.01	0.01	0.04	0.04	0.03
8	0.02	0.03	0.03	0.02	0.02	0.03	0.01	0.01	0.01	0.05	0.04	0.04
9	0.01	0.01	0.01	0.00	0.00	0.01	0.00	0.00	0.00	0.02	0.02	0.02
10	0.05	0.06	0.09	0.06	0.06	0.10	0.03	0.04	0.06	0.13	0.13	0.13
11	0.01	0.00	0.01	0.00	0.00	0.00	0.00	0.00	0.00	0.02	0.02	0.02
12	0.01	0.01	0.02	0.01	0.01	0.02	0.00	0.00	0.00	0.03	0.03	0.03
13	0.00	0.00	0.01	0.00	0.00	0.00	0.00	0.00	0.00	0.01	0.01	0.01
14	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.06	0.08	0.06
15	0.07	0.09	0.15	0.01	0.01	0.01	0.04	0.05	0.08	0.20	0.20	0.30
16	n.a.	n.a.	n.a.	0.04	0.05	0.11	n.a.	n.a.	n.a.	n.a.	n.a.	n.a.
Mean (%)	4.29	5.00	6.40	4.02	4.65	6.43	1.77	2.40	3.19	7.50	7.97	8.82

Notes: This table reports distributions of savings rates. $f_1(s^*)$ is the savings rate distribution of eligible employees at December 31, 1997, whose tenure is between three and five years. $f_2(s^*)$ is the savings rate distribution of eligible employees at December 31, 1997, whose tenure is between five and seven years. $f_3(s^*)$ is the distribution of optimal savings rates based on predicted values from an ordered logit regression of savings rate on age, gender, pay, and tenure. Predicted values are calculated using thirty years of tenure instead of actual tenure. n.a. = not applicable.

different test companies. The results of these minimizations are reported in table 2.3.

Table 2.3 documents six findings. First, the analysis reveals a high degree of heterogeneity in policy recommendations. The optimal default ranges from 0 percent to 15 percent. Moreover, even within a single firm there exists a large degree of variation in optimal defaults (e.g., Finance). Second, the range of variation in optimal defaults is twice as large as the range of average optimal savings rates. Third, the optimal default calculation is extremely sensitive to distributional assumptions on s^*. To see this, fix $\kappa = 100$ and read across the columns. The defaults show substantial variation arising from very small (within-company) differences in f_1, f_2, and f_3 (see table 2.2). Fourth, as κ gets large, much of the variation in optimal defaults is driven by the mode effect. For $\kappa = 1{,}000$, five out of twelve of the optimal defaults are equal to the modal optimal savings rate. Fifth, the optimal defaults vary in a sensible way with the underlying firm-specific attributes. Firms whose employees have a high motive to save turn out to have higher optimal defaults than firms whose employees have a low motive to save. For example, the employees in Food have a defined benefit plan and a low average salary (i.e., a high average Social Security income replacement rate) and hence very low optimal defaults (0 percent to 3 percent). By contrast, the employees in Finance have no defined benefit plan, a high average salary, and a median optimal default of 14 percent. Sixth, and finally, the optimal defaults tend to cluster in one of three regions: close to 0 percent, close to the match threshold (6 percent for Health and Office), or close to the maximum contribution rate allowed under the plan.

2.3 Discussion and Conclusion

This paper has presented a model of 401(k) enrollment. The model includes four components: costs of opting out of a default, an option value of waiting to incur those costs, procrastination in opting out of a default, and heterogeneity in optimal savings rates.

One should also consider other important psychological and economic issues when picking socially optimal defaults. First, some employees may interpret defaults as implicit advice, an issue that does not arise in the current model since each employee is assumed to know her true optimal savings rate.[8] Second, defaults may be particularly sticky because of loss aversion.[9] If the default is perceived to be a reference point, then deviations from that reference point may be psychologically aversive, since the resulting "gains" from the deviation (e.g., higher current consumption) are only

8. Employees may treat a zero default as weaker implicit advice than a nonzero default.

9. See Kahneman and Tversky (1979), Thaler (1980), and Samuelson and Zeckhauser (1988) for a discussion of loss aversion and status quo bias.

Table 2.3 **Optimal default savings rates**

	Health company			Office company			Food company			Financial company		
κ(%)	$f_1(s^*)$	$f_2(s^*)$	$f_3(s^*)$	$f_1(s^*)$	$f_2(s^*)$	$f_3(s^*)$	$f_1(s^*)$	$f_2(s^*)$	$f_3(s^*)$	$f_1(s^*)$	$f_2(s^*)$	$f_3(s^*)$
10	4	5	6	4	5	6	2	2	3	7	8	9
100	2	2	14	2	2	5	1	1	2	2	14	14
1,000	0	0	15	6	6	6	0	0	0	15	15	15
Mean (%)	4.29	5.00	6.40	4.02	4.65	6.43	1.77	2.40	3.19	7.50	7.97	8.82
Mode (%)	0.00	6.00	6.00	0.00	6.00	6.00	0.00	0.00	0.00	0.00	15.00	15.00

Notes: This table shows the optimal savings rate for four different firms. Food company and Financial company have no employer match in their plans. See table 2.2 notes for explanations of variables. Predicted values are calculated using thirty years of tenure instead of actual tenure.

weighted half as much as the resulting "losses" (e.g., lower saving). Third, if households do not know how to think about the future or are overoptimistic about future income, they may undervalue savings. In such a world, it may be optimal to pick a high default savings rate, even if households eventually move away from it. Fourth, households may know the optimal savings rate but not appreciate how important it is to implement it, increasing action delays. Fifth, choosing a long-run savings rate that is 1 percentage point too low is more costly than choosing a long-run savings rate that is 1 percentage point too high (since retirement is short relative to working life and the utility function generates a precautionary savings motive[10]), suggesting a desirable upward shading of optimal defaults. Sixth, optimal savings rates are not constant over time (as we assume) but instead are likely to trend up slowly with working age. Seventh, the firm may wish to pick an optimal default that weights some employees more heavily than others. For example, it may be sensible to calculate optimal defaults that overweight the interests of employees that are likely to have a long duration of employment at the firm and underweight employees that are likely to separate relatively quickly. Future work should extend our theoretical framework by incorporating many of these additional considerations.

Future work should also explore the empirical implications of our model. The model makes quantitative predictions about the timing of savings rate changes. Employees who change their savings rate soon after they are hired should select larger changes than employees who change their savings rate long after they are hired. This is because employees who are willing to wait a long time for a low-cost opportunity to opt out of the default are likely to have little to gain from doing so. The model also predicts that average savings rates will not necessarily increase monotonically with the default savings rate. As the default savings rate rises, procrastination effects can strengthen, leading more agents to delay selecting an even higher savings rate. Such perverse effects have already been observed in the data (Madrian and Shea 2001; Choi et al. 2004).

Finally, the model suggests one important generalization that we are currently exploring (Choi et al. 2002a). If it is occasionally optimal to select "bad" defaults—that is, defaults that are not close to one's optimum saving rate—then it may be optimal to pick defaults that are so bad that all consumers feel compelled to immediately opt out of them. Such a setup is equivalent in practice to something that we call "active decision," a regime that forces new employees to pick their own savings rate early in their tenure at the company without the benefit of a fallback default. In a world with significant procrastination, such active decision regimes are sometimes the best "defaults" of all.

10. Precautionary savings effects arise when $u''' > 0$, a common assumption in applied economic models.

References

Akerlof, George A. 1991. Procrastination and obedience. *American Economic Review Papers and Proceedings* 81 (2): 1–19.
Choi, James J., David Laibson, Brigitte Madrian, and Andrew Metrick. 2002a. Active decisions: A natural experiment in saving. Harvard University, Department of Economics. Mimeograph.
———. 2002b. Defined contribution pensions: Plan rules, participant decisions, and the path of least resistance. In *Tax policy and the economy,* ed. James M. Poterba, 67–113. Cambridge, MA: MIT Press.
———. 2003. Optimal defaults. *American Economic Review Papers and Proceedings* 93 (2): 180–85.
———. 2004. For better or for worse: Default effects and 401(k) savings behavior. In *Perspectives in the economics of aging,* ed. David Wise, 81–121. Chicago: University of Chicago Press.
Kahneman, Daniel, and Amos Tversky. 1979. Prospect theory: An analysis of decision under risk. *Econometrica* 47 (2): 263–92.
Laibson, David. 1997. Golden eggs and hyperbolic discounting. *Quarterly Journal of Economics* 112 (2): 443–77.
Laibson, David, Andrea Repetto, and Jeremy Tobacman. 2003. Instant gratification over the lifecycle. Harvard University, Department of Economics. Mimeograph.
Madrian, Brigitte C., and Dennis Shea. 2001. The power of suggestion: Inertia in 401(k) participation and savings behavior. *Quarterly Journal of Economics* 116 (4): 1149–87.
O'Donoghue, Ted, and Matthew Rabin. 1999. Doing it now or later. *American Economic Review* 89 (1): 103–24.
Samuelson, William, and Richard Zeckhauser. 1988. Status quo bias in decision making. *Journal of Risk and Uncertainty* 1 (1): 7–59.
Thaler, Richard H. 1980. Toward a positive theory of consumer choice. *Journal of Economic Behavior and Organization* 1 (1): 39–60.

Comment Antonio Rangel

The Question

A large number of companies and governments have introduced savings plans into the workplace. In a typical plan, individuals are allowed to contribute a percentage of their wages in exchange for a financial benefit such as a subsidized rate of return or a tax deduction. Since most plans offer significant financial advantages, one would expect participation rates to be high. Surprisingly, this is not the case. A series of recent papers have shown that a large fraction of employees either take too long to sign up for the plan (if they sign up at all) or fail to reoptimize their choices as their finan-

Antonio Rangel is an assistant professor of economics at Stanford University and a faculty research fellow of the National Bureau of Economic Research.

cial circumstances change. Subsequent research has shown that features of the savings plans that *have no effect on budget constraints,* such as the choice of defaults (Madrian and Shea 2001) or the timing and framing of questions (Benartzi and Thaler 2004), can have a sizable impact on the number and size of contributions.

These findings provide the motivation for the behavioral public economics question studied in this paper: How should a benevolent planner (e.g., the firm's benefit office) select defaults to minimize the mistakes made by individual workers?

Background

In order to understand the contribution and scope of this paper it is useful to start with a brief review of some of the psychological mechanisms have been proposed to explain the puzzling behaviors just described.

Transaction or decision-making costs (TCs). Deciding how much to contribute takes time and effort. Since the opportunity cost of these resources fluctuates with time, individuals are likely to wait for a period when TCs are low and to favor investment strategies that require little reoptimizing.

Procrastination. A sizable literature has shown (see, for example, Loewenstein, Read, and Baumeister 2003) that individuals tend to suboptimally postpone decisions that look like investments in the sense that they are costly in the present and generate benefits only in the future. This mechanism introduces an additional reason why employees may delay signing up for a savings plan, but only in cases where the financial costs of procrastination are not too large.

Imperfect attention. Individuals face a large number of decisions and routinely ignore most of them. In particular, unsophisticated decision makers may not think about savings unless appropriately cued to do so (Bernheim and Rangel 2003). Under this mechanism, workers postpone any choices related to the savings plan until they are exposed to a cue (such as an advertisement, a benefits fair, or a conversation with a family member) that helps them to focus their attention.

Fear of making mistakes and loss aversion. The psychological literature has shown that individuals postpone decisions when they are not sure about the right course of action. Some researchers have suggested that these fears are particularly paralyzing when poor choices can lead to financial losses.

Undersaving. All of the previous mechanisms provide reasons why workers may waste time in signing for the savings plan, but they cannot explain the low contribution rates. The mechanisms that have been proposed to explain that part of the puzzle include imperfections in internalizing the future benefits of savings, mistakes in calculating the amount of resources that are needed in retirement, and overoptimistic beliefs about investment returns.

It is important to emphasize that, with the exception of TCs, all of the mechanisms described in this section lead to systematic mistakes in decision making. A key assumption in this paper is that we can model the underlying psychological mechanisms at work and thus predict how mistakes change with behavioral features of the savings plan such as defaults.

A Sketch of the Model

The paper studies the role of default in a simple, elegant, and stationary model. Individuals are infinitely lived and need to choose how much they want to contribute to the savings plan every period. Once the choice is made, the contribution rate is fixed forever. Individuals differ on their optimal per-period savings rate s^*. Workers are required to contribute the default amount s^D unless they have explicitly signed up for a different rate. They incur a per-period quadratic *loss* in periods where their savings rate differs from s^* and do not discount future losses. Making decisions is costly. Transaction costs fluctuate stochastically: with probability μ the cost equals 1, with probability $1 - \mu$ the cost is zero.

Decision makers have a very special form of naive hyperbolic discounting. They always choose their optimal savings rate s^* when they show up to sign up for the plan. However, they may procrastinate in making that decision. More formally, the authors assume that individuals mistakenly overdiscount all future utility flows by $\beta \in (0,1)$ when deciding when to sign for the plan. Their naïveté also leads them to incorrectly believe that they will not make similar mistakes in the future.

We can now provide a more precise description of the question studied in the paper: How should a benevolent planner set up defaults when (a) there are time-varying TCs and (b) people procrastinate in making saving choices (but then act optimally)? Psychological mechanisms such as imperfect attention or loss aversion are not taken into account.

Contribution and Intuition

The paper develops the following three nice insights.

Insight 1: In the presence of decision-making costs, the choice of default matters even without procrastination. The intuition is straightforward. Let $L = (s^D - s^*)^2$ denote the per-period losses incurred by an individual who has yet to sign up for the plan. Straightforward computations show that individuals with $L > 1 - \mu$ sign immediately, whereas the rest do so in the first time TCs are zero. Figure 2C.1 plots the expected lifetime total losses (including decision-making costs) for individuals with different s^*. An individual with $s^* = s^D$ never signs up for the plan and thus never experiences losses or pays TCs. Individuals with an optimal contribution rate that is sufficiently different from the default sign up immediately. This group experiences a lifetime loss equal to the expected TCs in period 1. Finally, individuals with optimal contribution rates between s^L and s^H wait to sign up

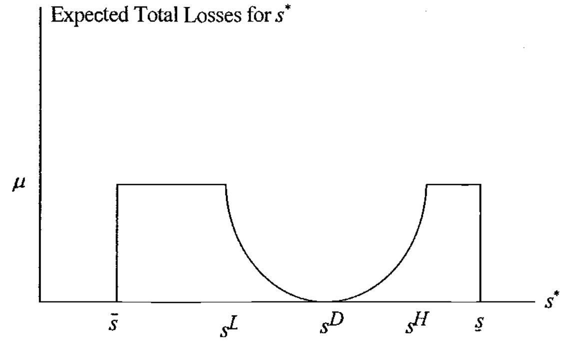

Fig. 2C.1 Losses without procrastination

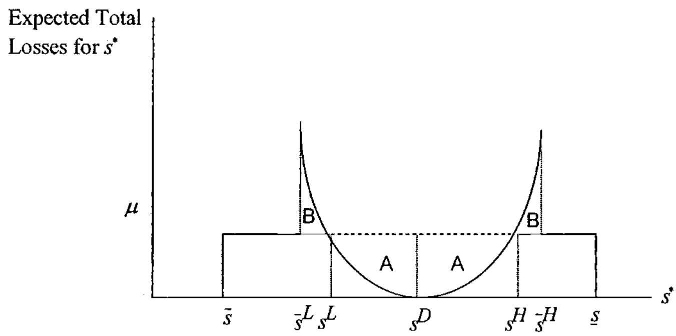

Fig. 2C.2 Losses with procrastination

for the plan until TCs are zero, and thus they experience an expected loss that increases with the square of $|s^* - s^D|$.

As can easily be seen in this picture, the optimal default depends on the distribution of individual preferences. If optimal saving rates are uniformly distributed between $\underline{s}$ and $\bar{s}$, then any default that is sufficiently far away from the corners is optimal. By contrast, if the distribution is a truncated normal centered around $(\underline{s} + \bar{s})/2$, the optimal default lies in the middle.

Insight 2: If there are TCs, procrastination, and sufficient heterogeneity, extreme defaults can be desirable. The intuition for this result is also simple. Straightforward computations show that the introduction of naive hyperbolic discounting induces a mistake for individuals with $L \in (1 - \mu, [1/\beta] - \mu)$: they should sign up immediately for the plan, but instead they procrastinate and wait until TCs are zero. The rest of the workforce behaves as before. This leads to the pattern of losses depicted in figure 2C.2. A comparison with figure 2C.1 illustrates the role that procrastination plays in the results. Individuals with optimal savings rates between $(\bar{s}^L, s^L)$ and $(s^H, \bar{s}^H)$

wait too long to sign up and thus experience additional losses. The number of individuals making mistakes and the size of their losses increase with the strength of the procrastination.

The potential attractiveness of extreme defaults follows immediately. Suppose that individuals are homogeneously distributed. If area B is larger than area A, then the optimal default is to set a default that is extremely high or low, which induces everyone to sign up in the first period. This increases social TCs by A but reduces the size of mistakes by an even larger amount.

Insight 3: In a large class of environments the optimal default is better than forcing all employees to make a choice. The intuition for this point follows immediately from the previous discussion since forcing everyone to make a choice is equivalent to picking a default that is sufficiently unattractive for everyone.

Other Issues to be Considered when Choosing Defaults

I conclude this discussion by developing some conjectures about what happens to the optimal default policy when additional psychological forces are at work.

Undersaving. Suppose that individuals choose suboptimally low contribution rates when they sign up for the plan. This would be predicted, among others, by a model where individuals exhibit hyperbolic discounting in all dimensions. Conjecture: When the undersaving mistakes are large enough, the optimal default looks like a mandatory contribution rate: TCs are increased to the point where all individuals stay with the default, and the default targets the average worker.

Imperfect attention. Suppose that there are no TCs but that individuals only think about savings in any given period with some probability p that is increasing in the size of the loss L. Conjecture: In this case the optimal default is close to average optimal savings rate and favors groups which, perhaps due to a lack of sophistication or education, exhibit lower probabilities of thinking about savings.

Inertia and Ignorance. Suppose that the presence of a default leads individuals to make a mistake when choosing their contribution rate. For example, they might mistakenly assume that the default is the right contribution rate for them. In that case forcing individuals to make a choice at time 1 might dominate the class of default policies studied in this paper (see Choi, Madrian, and Metrick 2002 for a related discussion).

References

Benartzi, Shlomo, and Richard Thaler. 2004. Save more tomorrow: Using behavioral economics to increase employee savings. *Journal of Political Economy* 112 (1): 164–87.

Bernheim, Douglas, and Antonio Rangel. 2003. Emotions, cognition, and savings: Theory and policy. Stanford University, Department of Economics. Manuscript.
Choi, James J., Brigitte C. Madrian, and Andrew Metrick. 2002. Defined contribution pensions: Plan rules, participant decisions, and the path of least resistance. In *Tax policy and the economy,* ed. James M. Poterba, 67–113. Cambridge: MIT Press.
Loewenstein, George, Daniel Read, and Roy F. Baumeister. 2003. *Time and decision: Economic and psychological perspectives on intertemporal choice.* Russell Sage Foundation.
Madrian, Brigitte, and Dennis Shea. 2001. The power of suggestion: Inertia in 401(k) participation and savings behavior. *Quarterly Journal of Economics* 116 (4): 1149–87.